

# Overview: Cloud Datacenters II

Hakim Weatherspoon

Associate Professor, Dept of Computer Science

CS 5413: High Performance Systems and Networking


January 30, 2017

# Background: The Internet




- How do we get bits into and out of datacenters?

# Background: The Internet

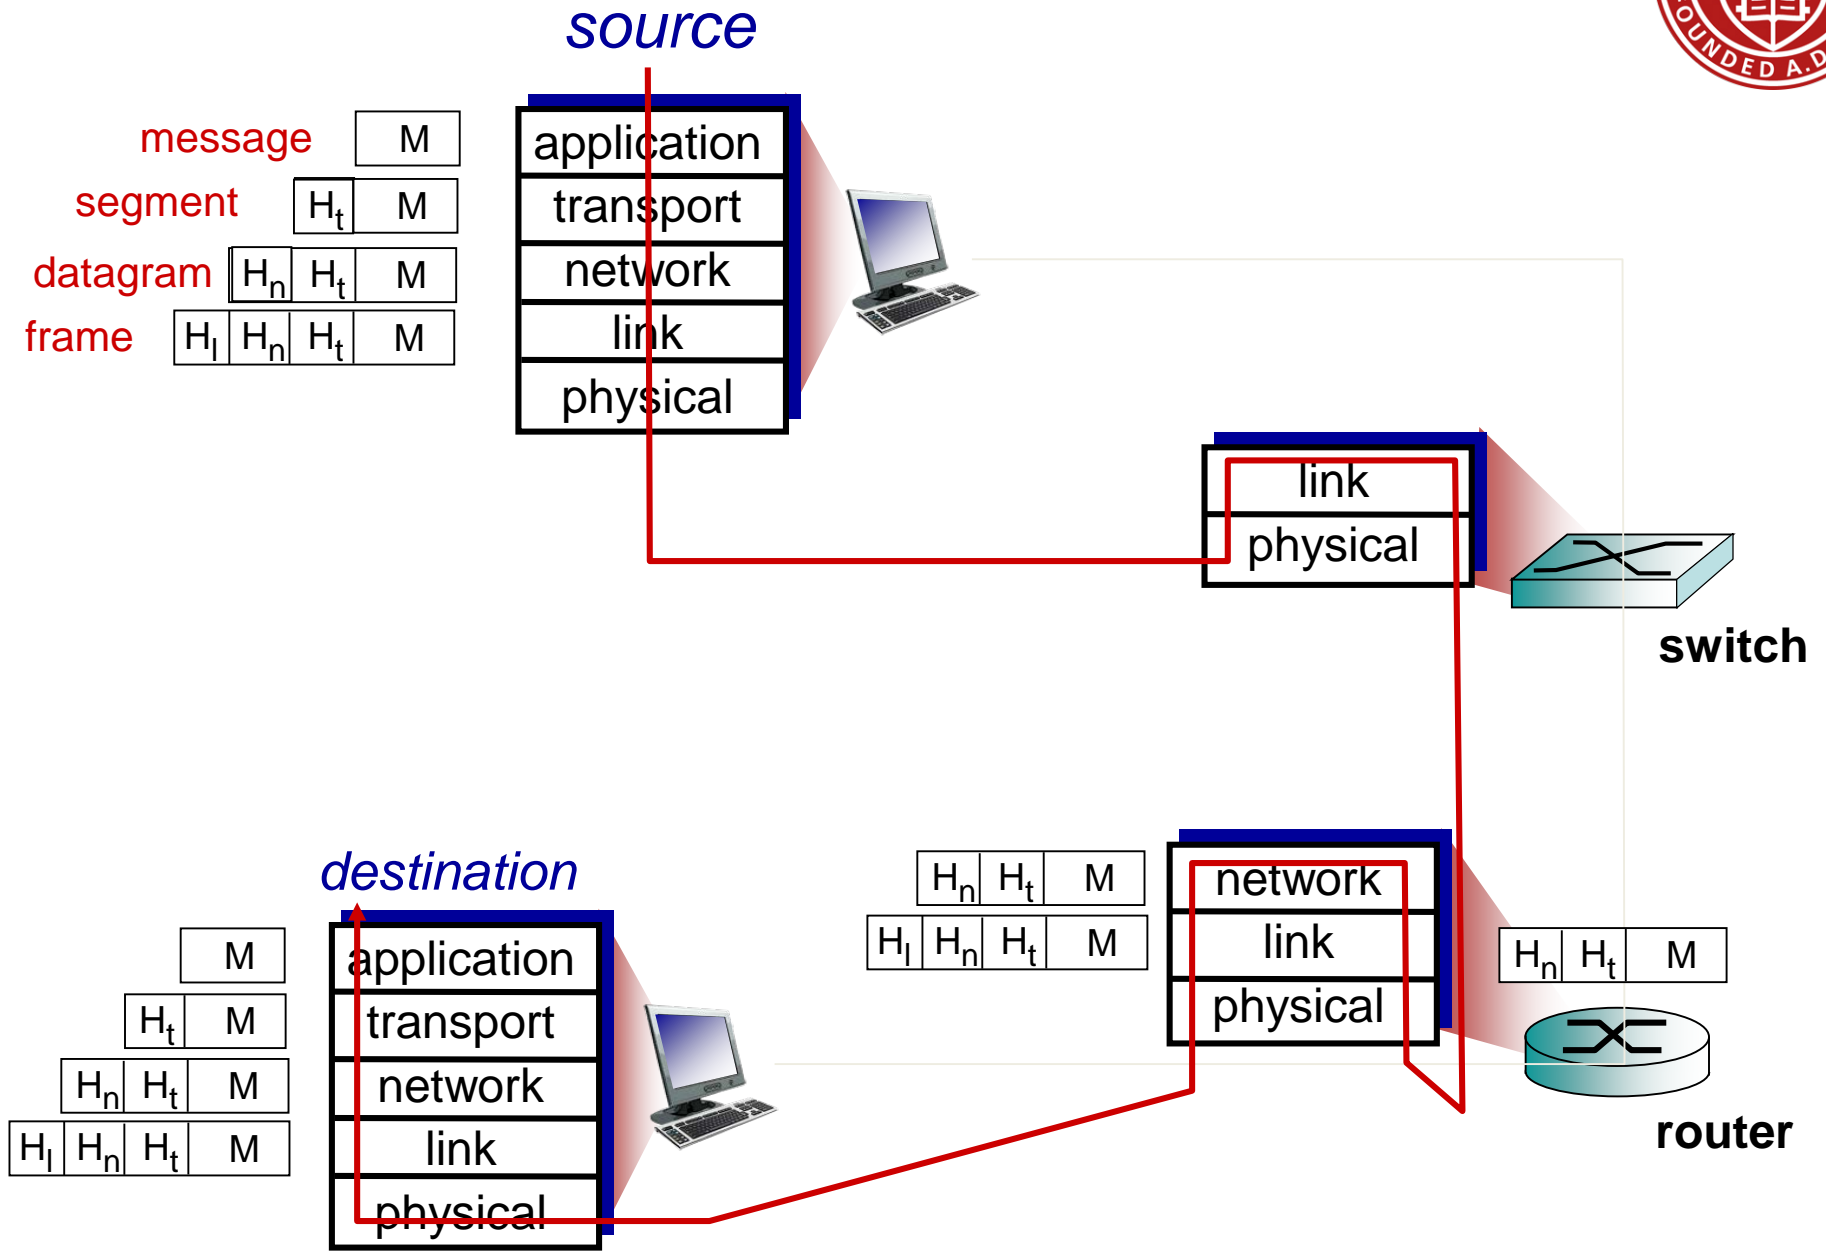



## Internet Protocol / Internet Protocol Stack

- *application*: supporting network applications
  - FTP, SMTP, HTTP
- *transport*: process-process data transfer
  - TCP, UDP
- *network*: routing of datagrams from source to destination
  - IP, routing protocols
- *link*: data transfer between neighboring network elements
  - Ethernet, 802.111 (WiFi), PPP
- *physical*: bits “on the wire”



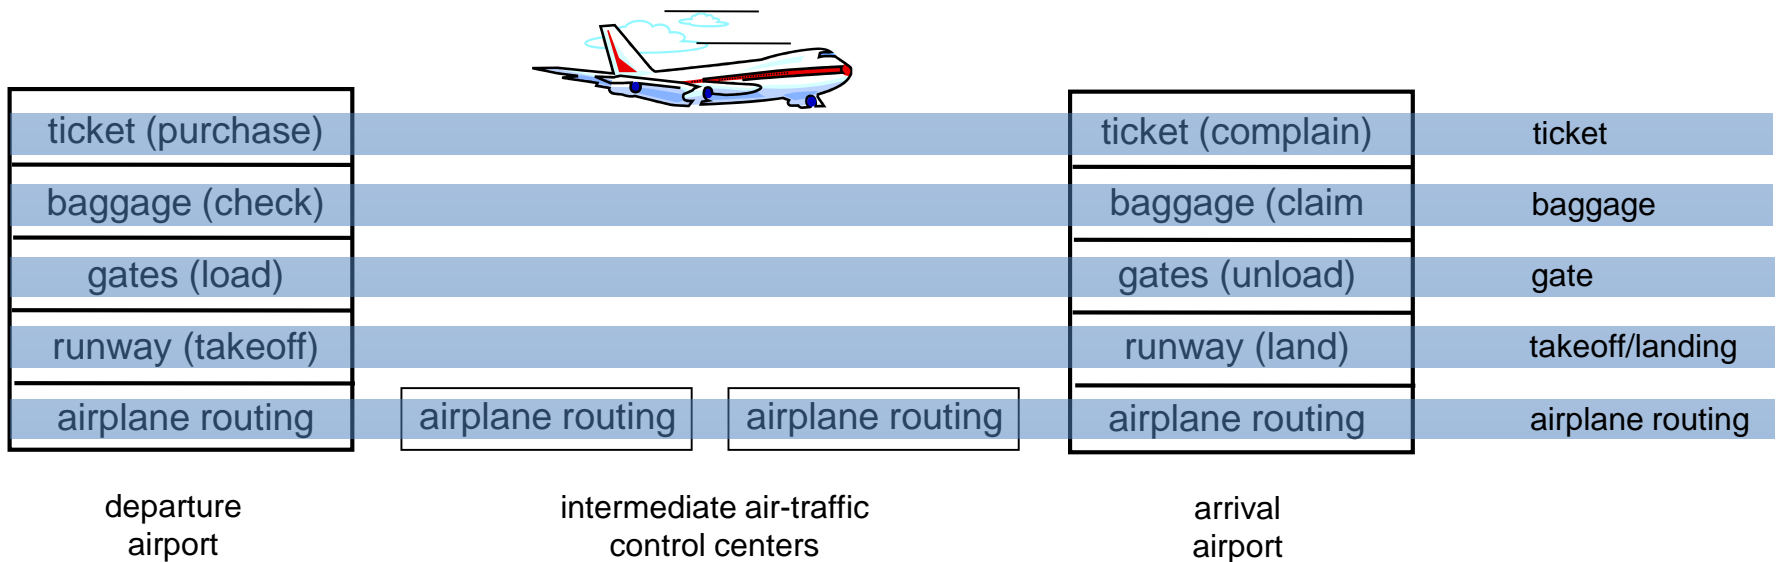
# Background: The Internet



# Network Protocol “Layers”




Network Protocol “Layers” similar to traveling protocol



*layers:* each layer implements a service


- via its own internal-layer actions
- relying on services provided by layer below

# Tech Titans Building Boom



- What does it take to build a million server datacenter?

# Tech Titans Building Boom



- What does it take to build a million server datacenter?
- Challenges
  - Readily available (fiber-optic) networking
  - Abundant water
  - Inexpensive electricity
    - How much electricity?
    - 200W per server \* 1M servers = 200MW!
    - Equivalent to 200k houses!
  - Management (e.g. installation, failures)
  - Environmental impact


# Tech Titans Building Boom



- What does it take to build a million server datacenter?
- Challenges
  - Readily available (fiber-optic) networking
  - Abundant water
  - Inexpensive electricity
    - How much electricity?
    - 200W per server \* 1M servers = 200MW!
    - Equivalent to 200k houses!
  - Management (e.g. installation, failures)
  - Environmental impact
- Prior state of the art, dot-com era of 1990's to 2000's
  - 1k to 2k servers -> 1MW to 2MW
  - Setup and management was fairly manual



# Tech Titans Building Boom



- What does it take to build a million server datacenter?
- Locations (power/cooling/water)
  - Washington, N.C., S.C., Iowa, Oklahoma,....,Siberia!




Titan tech boom, randy katz, 2008

# Tech Titans Building Boom




- What does it take to build a million server datacenter?
- Server Utilization
  - 40x 200W pizza boxes
  - CPUs are 60% of power
  - 8 to 16kW per rack
  - 0.5kW/m<sup>2</sup>
  - Air cooling



- Google/Microsoft
  - Better power mgmt.
    - . Avg instead of peak
  - Better power supplies
    - voltage regulators, fans
  - Remove GPU
  - Water cooling


# Tech Titans Building Boom




- What does it take to build a million server datacenter?
- Containers (server, power, cooling efficiency)
  - 2500 to 3000 servers, instead of 40 to 80
  - Power and cooling efficiency
  - Power density,  $16\text{kW}/\text{m}^2$  instead of  $0.5\text{ kW}/\text{m}^2$



# Tech Titans Building Boom




- What does it take to build a million server datacenter?




Titan tech boom, randy katz, 2008

# Tech Titans Building Boom




- What does it take to build a million server datacenter?
  - Power efficiency
  - Cooling efficiency
  - Server efficiency
    - Power proportionality
    - utilization
  - Power density
    - 0.5 kW/m<sup>2</sup> – raised floor datacenter
    - 16 kW/m<sup>2</sup> – containerized datacenter
  - Management/failure
    - Software masked failures
    - containerization

# Tech Titans Building Boom




- Power efficiency
  - Tune power supply for average, not peak
  - Voltage regulators
  - Remove unnecessary components
- Cooling efficiency
  - HP “smart cooling”
  - Air-side economization
  - Containers

# Tech Titans Building Boom



- PUE
  - Total power consumption / total power used by consumers
- Results
  - Typical enterprise DC
    - 2007 – 2
    - 2011 – 1.7 (with optimizations may reach 1.3)
  - Google DCs
    - Avg – 1.21
    - Best – 1.15
  - Microsoft
    - Chicago – 1.22


# Tech Titans Building Boom



- Virtualization
  - DCs run at 15% of their capacity without virtualization
  - DCs run at 80% with virtualization
- Other SW tools
  - Power usage control
  - Shared distributed data
  - Handle software failures




# Perspective



- To build large and efficient datacenters
  - Better power efficiency
  - Better cooling efficiency
  - Specialized systems for datacenters

# *Before* Next time



- Finish Lab0 by Tuesday
- Fill out survey to help form groups
- Create a project group
  - Start asking questions about possible projects
- Check website for updated schedule